



Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

Fall 2015 HEPiX Workshop – Fermilab Site Report

Rennie Scott

12 October 2015

Fermilab - Introduction

- Rennie Scott
- Fermilab – 8 years
- Previously NCSA, and SGI
- Manager - Scientific Linux and Architecture Management
 - Experiment Online infrastructure
 - Operations Control Rooms
 - Scientific Workstations
 - Scientific Linux

Fermilab - Introduction

- US DOE – Laboratory
 - Located approximately 45 miles outside downtown Chicago
 - Established – 46 years ago
 - Approximately 1,750 employees

Computing Facilities on the Fermilab site

3 main computer room locations

- **Feynman Computing Center (FCC)**

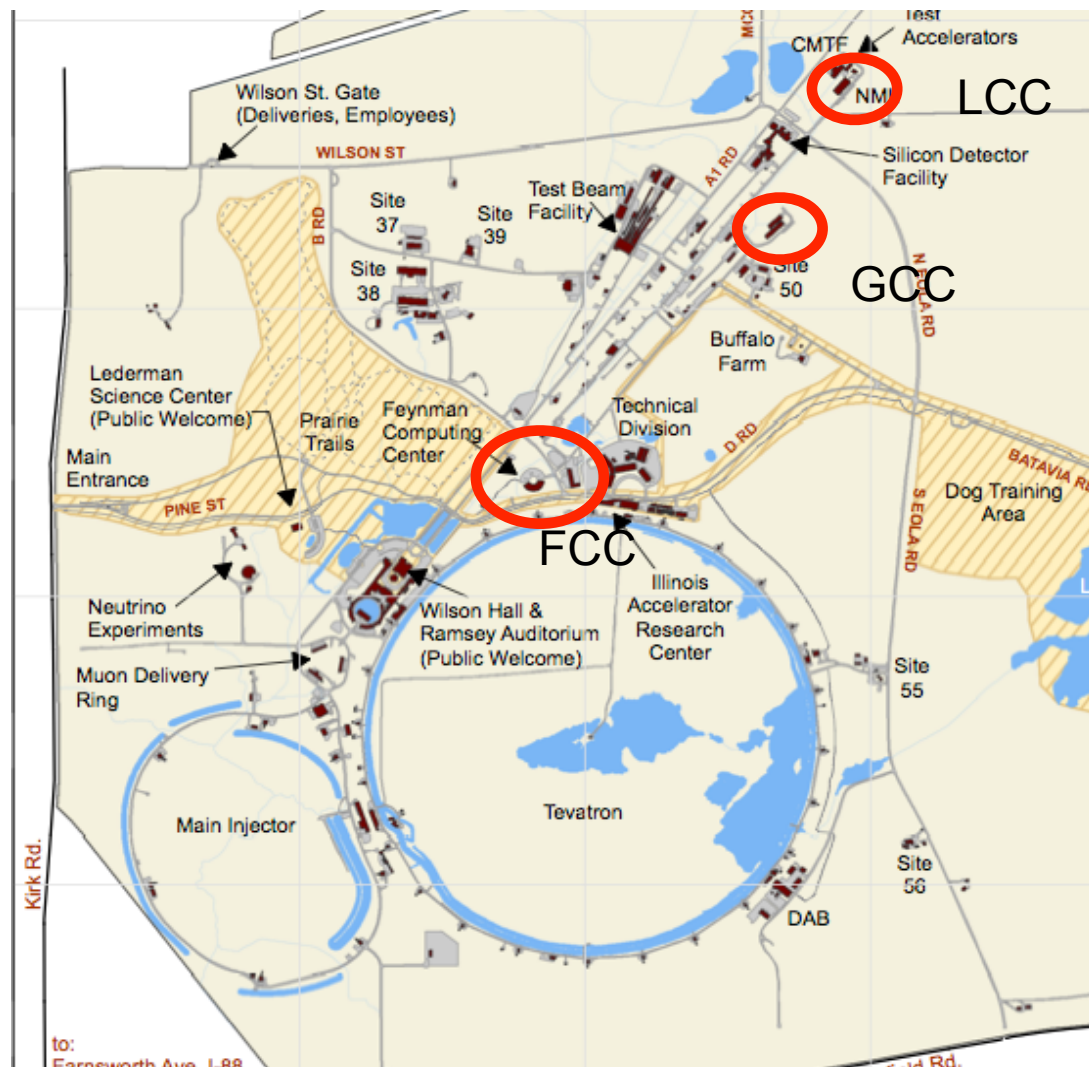
- 2 rooms with 0.75 MW nominal cooling and electrical power each
 - UPS with independent generator backup
- Hosts power critical services
 - central services (Mail, web servers, etc.) and disk servers

- **Grid Computing Center (GCC)**

- 3 rooms with 0.90 MW nominal power each
 - UPS with taps for external generators (no permanent generator)
- Hosts CPUs and Tape libraries, UPS sustains power during power outages till systems can be powered down in a controlled way

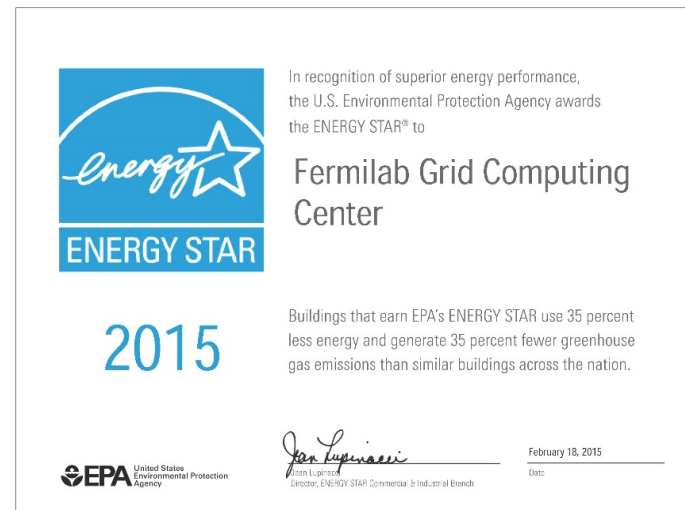
- **Lattice Computing Center (LCC)**

- 1 room with 0.47 MW nominal power
 - No UPS and no generator
- Houses CPU and GPU clusters running less power-critical applications



Computing Facilities

- Data Center reliability - one of the best years – so far
- Retired 1079 sqft of computer room space
 - Room 108 at LCC in July, 2015; repurposed the space
- Energy Star for the Grid Computing Center
 - Five consecutive years



ISO 20K / ITIL Service Management

- Service Document Streamlining to one consolidated document covering descriptions, SLA, Capacity, Continuity, and Availability. Displayed in ServiceNow going forward.
- ISO20K Recertification Audit – October 26-29, 2015
- Onboarding all Scientific Service Areas and most offerings by end of 2015.
- Technical Scope of Work (TSW) for experiments going well
 - Complete: MicroBooNE, MINOS+, NOvA, SeaQuest
 - Out for signature: LArIAT
 - Out for comment: DUNE 35T, MINERvA

Authentication/Identity Management

- KCA Server EOL Announcement
 - Kerberos Authenticated Certificate Authority
 - Service to be shutdown on September 30th, 2016
 - Application reliant on deprecated code
 - Web services moving to Single Sign On
 - Grid services evaluating CiLogon certificates

Security updates


- Website security – reducing attack surface
 - Move to central web services
 - Identifying content owners
 - Response to DOE request to reduce unnecessary websites
- Migration of scanner farm to new code
 - Integration with application team servers
 - Faster speed
 - Old servers were hard to maintain and were slower

Web services

- WordPress Service Rollout
- Organization websites moved to Hitachi NAS System
- PLONE Retirement – Oct 30, 2015

Distributed Computing

- 33 Fermilab experiment groups submitting batch jobs to one condor cluster using a common job submission utility developed at Fermilab called jobsub
- ~26,000 cores on site (15k on CMS LPC and T1 farms, 11k for the rest).
- Experiments successful getting opportunistic cpu cycles from OSG sites using glideinwms



Distributed Computing (cont)

- General purpose farm (GPgrid) has always allowed access to disk space on a central Bluearc NFS server but we will remove that capability
- Moved to distributing production analysis code via CVMFS.
- Central NAS server will still be used for test versions of analysis code for quick turnaround.
- All user data input and output from batch jobs is moving via grid protocols.
 - Persistent dCache is scratch space for this data.

Distributed Computing (cont)

- Using Grafana for monitoring of batch jobs. Very easy to import time series data and very easy to create custom dashboards with the data.



LQCD update

- Expanded existing “pi” CPU-based cluster from 214 to 314 nodes:
 - Expansion consists of 1,600 Intel Xeon cores delivering 6TF**.
 - Allows calculations related to the Muon g-2 experiment to increase job sizes
- Upgraded a critical *XFS-over-NFS* server based on spinning disks with *ZFS-over-NFS* based on a mix of Solid State and spinning disks.
 - Fast write cache using SSD drives,
 - End-to-end data integrity, snapshots and integrated data compression,
 - Serving 1,100 clients with a 1.5x speedup on read operations.
- Upgrading our existing 1PB Lustre version 1.8 based on the ldiskfs file-system to a stable recent release. Since in-place upgrade is not possible a data migration plan is underway which involves the following steps:
 - Deployed into production a 200TB Lustre version 2.5 based on the ZFS file-system.
 - Migration of data and storage hardware from Lustre 1.8 to 2.5 under way and plan to complete by end of CY2015. Migrating 89M files with 800TB aggregate data.
 - Both Lusters' are based on Infiniband QDR with 1,100 clients.

** based on LQCD benchmarks

DUNE/LBNF

- DUNE/LBNF pass DOE CD-1 Refresh Review
 - This is an important milestone and confirms that the conceptual design is sound and development can proceed.
- LBNF preparing for DOE CD-3a Review in December, to allow underground construction at Sanford Underground Research Facility to begin.
- DUNE 35 ton experiment
 - The detector is scheduled to turn on the week of Thanksgiving.

Muon g-2


- The magnet was successfully cooled and powered to full field.
- A 25% prototype of the DAQ has been built and tested at the full expected data rate for g-2.
- Front-end systems will consist of 12 systems with 24 GPUs (Tesla K40C 12GB) 240GB SSD
- Back-end 10 TB Raid

Mu2e

- Plan to move to ROOT6 once the art software stack has finished their migration.
- Making heavy use of non-Fermilab OSG sites.

MicroBooNE

- 170 ton Liquid Argon Time Projection Chamber (LArTPC) located along the Booster neutrino beam line.
- First images of cosmic rays recorded by the TPC!



Fermilab Test Beam Facility



- The goal of the Fermilab Test Beam Facility (FTBF) program is to provide flexible, equal and open access to test beams for all detector tests, with relatively low bureaucratic overhead and a guarantee of safety, coordination and oversight.
- We recently completed a Small board Project collaboration using an Intel Minnowboard is in production for 4 second beam AD signaling. It's running minimally modified version of Scientific Linux and managed by Puppet

HEPCLOUD Project

- The goal of the Fermilab HEP Cloud Facility Project is to extend the current Fermilab Computing Facility to transparently run on disparate resources including commercial and community clouds, grid federations, and HPC centers to meet peak demands or special requirements.
- Currently focusing on two cloud demonstrators:
 - CMS – Demonstrate Scale, 56K cores continuous for 1 month
 - NOvA – Demonstrate peak demands, smaller workflows, multiple times over the course of several months

Scientific Software (cont)

- art is an event-processing framework for particle physics experiments..
- Used by a growing number of neutrino and muon experiments, primarily at Fermilab but also elsewhere.
- Part of the underpinnings of both artdaq and LArSoft.
- This summer we put on a week-long course teaching the use of art and LArSoft, to 50 experimenters from around the world, who came to Fermilab to learn how to use art.

Scientific Software

- *artdaq* is a framework for developing software for data acquisition (DAQ) systems. Many DAQ functions (e.g. data transfers and event building) are provided by the framework, and users only need to provide the software that is specific to their experiment.
- It is used by a growing number of neutrino and dark matter experiments that are based at, or connected with, Fermilab.

Upcoming Fermilab HEPiX talks

- Natalia – Space usage monitoring for distributed heterogeneous data storage systems.
- Dmitry & Gerard– Scientific Data Storage at FNAL
- Bonnie ITIL Service Models in Detector DAQ Computing

Questions?